

2019 Minerva User Survey Comments and Response

2020-02-26

System Issues

- The hydra filesystem was extremely slow during the upgrade 2019; So far in 2020, there are no major issues with the filesystems.
 - We are sorry for the sub-optimal performance in hydra file system during the Minerva upgrade 2019. This issue has already been fixed by the PM on Dec. 8 2019. This latency issue with write operations on hydra is a GPFS v5.0.2 bug, which is fixed by reconfiguring a GPFS parameter.
 - We will keep monitoring the file system performance and make adjustments as needed in time in future.
- Too many maintenance disruptions. Please do not schedule maintenance during business hours and send multiple reminders.
 - In 2019, we had the complete refreshment of the Minerva cluster: retire all the old compute nodes; add ~18,000 new intel compute nodes; upgrade the GPFS to the latest version with bug fix and better feature; file migration between file systems; upgrade OS; upgrade all the IB network. Some of the operations will affect the production system and can only be performed in a maintenance period.
 - In 2020, our major task is adding the newly purchased storage DSS of 14 PB of raw storage and flash disk to the arion file system, which is already completed in May 2020. We have performed the two PMs needed for this task on Saturday. We are not expecting much PMs going forward, and we will always keep in mind to minimize the number of PMs. We will also try our best to schedule the PMs on Weekends and send out announcements 2 weeks ahead, with reminders followed.
- Please bring back the dedicated data transfer nodes and improve download speeds
 - We are working on the dedicated data transfer nodes. For now, we have no CPU time limit on minerva12(li03c02) and we have globus running on minerva11(li03c01), which you can use for data transfer. The maximum network speed is 10GB/s on login nodes. It is limited by the speed of the campus network and we can not speed up more than that. You will be competing with other users for the network traffic, and doing data transfer while using VPN is not recommended.
- Web server: Fix the web hosting thing please
 - We have set up the new web server on Minerva, and migrated the old web service, with all the file systems mounted.
- Scratch gets filled quickly. I suggest removing unused files more often

- We have implemented a per-user space limit on scratch, which is 15TB. And we have a policy running daily to purge the files every 14 days. A lot of users need those source/intermediate data for their pipeline processing, so it is not convenient to those users if we purge too often.

Job scheduler & Queues

- Sometimes, the LSF is down and some nodes do not function well, resulting in a very long queue and job submission issue; huge delays in executing pending jobs; Queuing time has been increasingly long since the Chimera shift in mid-2019
 - We updated LSF job scheduler with the latest fixes, resolving the bug that caused lots of jobs in pending state
 - Resource limit is lifted on the LSF server to reduce server crash
 - We set maximum pending jobs for each user as 20,000 to prevent users from submitting several (tens of thousands) short jobs (O(minutes)), which query LSF too often
- Sometimes too long waiting time. Sometimes I had difficulty getting computation resources because some individuals used up all of them.
 - We have added 78 extra BODE nodes in Nov. 2019 for NIH-funded projects, which helps reduce the waiting time.
 - We have fixed the two issues which caused a lot of job failure and delayed the processing of the batch jobs (slow I/O in hydra and permission denied by account manager).
 - Set limit on running jobs for heavy users
 - Adjust queueing parameters according to user jobs
 - We are going to purchase more new compute nodes with larger-shared memory (1.5 TB). Once it is in production (estimated in 2020 Q4), the waiting time of jobs will be further greatly reduced.
 - Run into maintenance
 - Added default check for users with immediate screen message
- LSF vs SLURM
 - We received a mixture of opinions on switching our LSF job scheduler to SLURM to save cost. So we will not switch to SLURM this year. We will revisit this with you users next year.
- Long waiting times on the gpu queue

- We are going to purchase more GPU nodes with 4*A100 each node. Once it is in production (estimated in 2020 Q4), the waiting time of jobs in GPU queues will be further greatly reduced.

TSM

- TSM is very frequently complains that the data is not available on the server and communication with hpchelp is needed. It slows the retrieval process by 1-2 days. Achieving large data sometimes have cache full issue
 - We currently have a total of 19.7PB data on TSM which sits on more than 12,000 tapes. Unfortunately, the current TSM library is not big enough to accommodate all these tapes. When a retrieve command is evoked, the operators have to fetch the tape and load manually into the library. The responding time depends but should be less than 1.5h. Once the tape is loaded, the data transfer starts at a constant rate of 150M/s. The data unavailable error is due to some error that is raised on the retrieve process and needs manual interference to fix. The cache full is due to lack of scratch tapes. Those errors are evitable but we will try to monitor the TSM more closely to avoid situations like this.
- TSM is really difficult to use with little instructions on your website and a lot of guess work to make it do its job.
 - Our detailed documentation on TSM is at <https://labs.icahn.mssm.edu/minervalab/archiving-data/>

Software/packages

- Having Snakemake available as a module would really help.
 - Snakemake is available via python packages. To access, use

```
$ ml python
$ which snakemake
/hpc/packages/minerva-centos7/py_packages/3.7/bin/snakemake
```
- It would be great to have notifications when modules removed or default version updated
 - We don't want to spam the users' email with notifications whenever there is a version change for every module, since there are over 1000 modules. But we will notify users when major packages have version changes such as R and python.
 - To check what version is available and is default, you can use the command "ml avail" or "ml spider".

- Remove old/outdated software
 - We are trying our best to clean up the unnecessary packages, but we cannot really remove all the old software, since many users still rely on those for reproduction and consistency of their publications.
- Will it be possible to have libre office installed so that the word/excel files can be seen in the chimera rather than we have to copy the files locally
 - We will try to install it or a similar light version on Minerva.
- Current setup works well but being able to use Jupyter notebooks and browser access would be handy
 - Jupyter Notebooks are available in Minerva. Please access through module system:
 - `$ ml python`
 - `$ which jupyter`
 - `$ /hpc/packages/minerva-centos7/py_packages/3.7/bin/jupyter`
 - If you want to access via browser, please check the procedure at <https://www.digitalocean.com/community/tutorials/how-to-install-run-connect-to-jupyter-notebook-on-remote-server>

Tickets

- When tickets are addressed by the staff they usually do a very good job of addressing the issue. However there are often significant delays before a ticket gets a response. I do not think it's an issue with the quality of the staff. There are simply not enough sys-admins to support the number of users currently active on the cluster. staff is overcommitted due to lack of manpower; I feel the employees are working very hard but there are delays due to sheer volume of requests; Sometimes I never got any response.
 - Yes, the low staff is challenging. This has caused slower than usual response time to help researchers on some specific topics. We have two vacant positions open for HPC admins. We are trying very hard to recruit HPC admins. We already have one candidate who is willing to join us.
 - We have about a ~25% increase in the number of tickets in 2019 due to the system upgrade and migration. We will try our best to reduce the time taken for ticket response this year.
- There were times that HPC staff had other priorities other than answering tickets, such that it took more than a week to hear a response.
 - Yes, sometimes we have other deadlines such as grant, report, invoice et al. We have tried our best to devote some time taking care of the urgent tickets during this period. We will have an additional system administrator starting in early August, which will relieve some of the burden on the

Computational Scientists. We also are continuing to look for high quality candidates to fill other open positions to further bolster our help desk staff.

Tutorial/Training

- Having more workshops on how to use Minerva
- Generally, I find that hpc support is incredibly helpful and fast. I have found, however, that for basic/introductory questions on using a cluster or working in bash at the command line, it is difficult to get answers to simpler questions. I think the tutorial sessions could include more basic information (or could have a very introductory session). Even a hello world for how to submit a job, check it's status, and confirm output would be helpful. Some information supporting the use of git on minerva would be helpful as well since sinai has a github enterprise suite that could be more openly advertised.
 - We offered two tutorial training in 2019, including both “Introduction to Minerva environment” and “Submitting jobs via LSF”. This year, we have offered another two training in March. We will continue to schedule tutorials this Fall.
 - The latest slide is at <https://labs.icaahn.mssm.edu/minervalab/training-archives/>.
 - If your research group would like to schedule an additional tutorial or tutorials focused on special HPC topics, you can reach out to our computational scientists (hpchelp@hpc.mssm.edu).

Documentation on Website

- The documentation is terrible and hard to navigate and lacks information. It is impossible to use it when you are a new user. A decent amount of it seems to be out of date and there is no obvious way to tell what is current and what isn't.
 - Documentation is clearly insufficient. We have started a project to revamp and update the documentation on Minerva's website. For the most recent updates and guides, please read/follow our announcement and Newsletter (subscribe to hpcusers@mssm.edu). We will update the documentation on our website in a timely manner to reflect the changes in our system.
- Information that would be helpful include:
 - where to store and run jobs
 - Architecture of a cluster computer (nodes vs cores vs tiles), size/speed of resources available, documentation on how to make scripts communicate

between cores, suggested best use practices (e.g. selfsched for batch job submission)

- Please check architecture of a cluster computer and resource at https://labs.icaahn.mssm.edu/minervalab/wp-content/uploads/sites/342/2020/03/Minerva_Introduction_-2020-03-11.pdf
- Please check our examples for batch job submission and selfsched at our latest slides at https://labs.icaahn.mssm.edu/minervalab/wp-content/uploads/sites/342/2020/03/Minerva_LSF_2020-03-18.pdf
- For Minerva Resources on Website, please check <https://labs.icaahn.mssm.edu/minervalab/resources/> .

Cloud vs Minerva

- Provide and improve the cloud services
 - We support singularity container on Minerva, which you can access by ml singularity. You can run docker/singularity image and modify an existing images with sandbox. More details at page 15 of slide https://labs.icaahn.mssm.edu/minervalab/wp-content/uploads/sites/342/2020/03/Minerva_Introduction_-2020-03-11.pdf or singularity quick start at <https://singularity.lbl.gov/quickstart>
- Adopt AWS cloud and their support service.
 - AWS cloud is more expensive and less efficient for the scale of our science; We estimated the costs using our actual compute core hours and storage used by researchers on Minerva for the preceding three years (~73% of the total available compute core hours) and 50% of the total available storage over that same period, as a minimal baseline. We found the cloud costs to be approximately three times higher than the cost of hosting the infrastructure ourselves. At our scale of computing and storage usage, the cloud is cost-prohibitive if remote cloud vendors are utilized.

Thank you for your nice words!

- Staff is always unbelievably wonderful.
- Whenever I need software help, the help team is very fast and responsive! I love it!
- Content was always good, and if I needed something HPC help was great at installing it.

- Staff responsiveness is great!
- Amazing. The team is doing a great job. Thank you for all your support.
- I appreciate kind and timely responses from hpchelp
- Staffs are extremely helpful and usually response to our questions very quickly and professionally
- HPC does a great job
- Staff are always very responsive and helpful when I contact the helpdesk, with very quick response times.
- Staff is wonderful, and they respond to tickets promptly.
- We are very happy with the service!!
- I am very happy with the service from the staff - they have been pushed very hard this year because of all of the upgrades, and the difficulties that entailed!
- Staff are very responsive.